

ARMY RESEARCH LABORATORY



The Human Side of Automation: Lessons for Air Defense Command and Control

by John K. Hawley, Anna L. Mares, and Cheryl A. Giannamico

ARL-TR-3468

March 2005

NOTICES

Disclaimers

The findings in this report are not to be construed as of an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

DESTRUCTION NOTICE—Destroy this report when it is no longer needed. Do not return it to the originator.

Army Research Laboratory

Aberdeen Proving Ground, MD 21005-5066

ARL-TR-3468

March 2005

The Human Side of Automation: Lessons for Air Defense Command and Control

John K. Hawley, Anna L. Mares, and Cheryl A. Giannamico
Human Research and Engineering Directorate, ARL

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY) March 2005	2. REPORT TYPE Final	3. DATES COVERED (From - To) May 1 to September 2004		
4. TITLE AND SUBTITLE The Human Side of Automation: Lessons for Air Defense Command and Control		5a. CONTRACT NUMBER		
		5b. GRANT NUMBER		
		5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S) John K. Hawley, Anna L. Mares, and Cheryl A. Giammanco		5d. PROJECT NUMBER		
		5e. TASK NUMBER		
		5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) U.S. Army Research Laboratory ATTN: AMSRD-ARL-HR-MB Human Research and Engineering Directorate Aberdeen Proving Ground, MD 21005-5066		8. PERFORMING ORGANIZATION REPORT NUMBER ARL-TR-3468		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army Research Laboratory Aberdeen Proving Ground, MD 21005-5066		10. SPONSOR/MONITOR'S ACRONYM(S)		
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited.				
13. SUPPLEMENTARY NOTES				
14. ABSTRACT One of the defining properties of the next generation of air defense command and control (C2) systems is an increasing reliance on automation. This paper provides an overview of human performance and training issues in automated air defense C2 systems. Its primary focus is the impact of automation on air defense C2 operators and the consequences of their role change from traditional operator to supervisors of automated processes. Topics that are discussed include (a) the changing role of operators under automation, (b) human performance and automation, (c) the operator's role as a supervisory controller, (d) problems with undisciplined automation, (e) developing effective automated systems, and (f) the impact of automation on training and staff development. The paper is intended as a primer on automation and human performance for commanders, concept developers, system designers, trainers, and other personnel involved with decision-making for the next generation of automated C2 systems.				
15. SUBJECT TERMS Automation, supervisory control				
16. SECURITY CLASSIFICATION OF: a. REPORT Unclassified		17. LIMITATION OF ABSTRACT UL	18. NUMBER OF PAGES 36	19a. NAME OF RESPONSIBLE PERSON John K. Hawley
b. ABSTRACT Unclassified	c. THIS PAGE Unclassified			19b. TELEPHONE NUMBER (Include area code) 410-278-5867

Standard Form 298 (Rev. 8/98)
Prescribed by ANSI Std. Z39.18

Contents

List of Figures	iv
List of Tables	iv
1. Introduction	1
1.1 Overview	1
1.2 The Problem	1
1.3 Air Defense Command and Control as Process Control	2
1.4 The Changing Role of Operators Under Automation	3
2. Automation and Process Control	4
2.1 An Overview of Automation.....	4
2.2 Human Performance and Automation.....	6
3. Human Supervisory Control	10
3.1 Traditional Versus Supervisory Control.....	11
3.2 The Human Operator's Role as a Supervisory Controller	12
4. The Path Forward	19
4.1 Problems with Undisciplined Automation	19
4.2 Developing Effective Automated Systems.....	23
4.3 Closing Perspectives: Beyond Human-Centered Design	24
5. References	27
Distribution List	30

List of Figures

Figure 1. A representative person-machine process control system.....	3
Figure 2. Basic supervisory control paradigm.....	11
Figure 3. Rasmussen's supervisory control taxonomy.....	13
Figure 4. Rasmussen's supervisory control taxonomy and human-system function allocation in an automated system.....	14
Figure 5. Functional and temporal nesting of generic supervisory control functions.	15
Figure 6. Operator interaction with a computer-based decision aid.	16
Figure 7. Potential combinations of human and computer control.....	18

List of Tables

Table 1. Levels of contemporary automation.	6
--	---

1. Introduction

1.1 Overview

One of the defining properties of the next generation of air defense command and control (C2) systems is an increasing reliance on automation. This paper provides an overview of human performance and training issues in automated air defense C2 systems. Its primary focus is the impact of automation on human operators and the consequences of their role change from traditional operator to supervisory controller. The paper is intended as a primer on automation and human performance for commanders, concept developers, system designers, trainers, and other personnel involved with decision-making for the next generation of automated C2 systems.

1.2 The Problem

Over the past forty years, the air defense C2 function has evolved from the relatively straightforward tasks of the gun-based era to the highly demanding and complex tasks required to optimize the employment of today's missile-based systems. Gun-based systems relied in large part on placing a "curtain of steel" to defeat an in-bound hostile threat. Target detection was based on visual and acoustic means. And engagement decision-making was simple: If a target threatened you or the assets in your area, you fired upon it. Furthermore, little effort was made to coordinate the actions of individual fire units.

Early developments in radar technology did not significantly alter the C2 process. While the ranges at which targets could be detected and the accuracy of determining the direction of approach was improved, the engagement itself was still largely an autonomous process. As far as the individual elements of C3I (Command, Control, Communications, and Intelligence) were concerned, the intelligence ("I") function was evident, and perhaps some communication (the third "C") was involved in the distribution of intelligence information. Command and control of air defense assets was, however, still limited principally to siting assets in positions where unit commanders determined that they could be most effective.

Air defense C2 as it exists today began to evolve in the mid-1950s. The emergence of a long-range bomber threat drove the development of missile-based air defense systems. The need to defeat weapons of mass destruction at long range became paramount. Moreover, the cost of missile-based weapons systems mandated a change in tactics. No longer could one afford the luxury of relying upon a steel curtain to defeat an air threat. The goal became the selection of the most appropriate air defense unit to fire one, or at most two, missiles to destroy each in-bound threat. Also, since both fighter interceptor aircraft and ground based missiles might be called upon to engage a threat, joint C2 of both of these types of assets became vital to effective threat engagement, precluding fratricide, and optimal use of limited resources.

The information technologies that have evolved over the past several decades have provided the building blocks for improved C2 systems. In the late 1950s, 1960s, and 1970s, the AN/TSQ-51 Missile Mentor and, later, the AN/TSQ-73 Missile Minder systems provided air defense commanders with an electronic means for managing air defense assets. Commands, status information, and intelligence data were transmitted and received by electronic data links, greatly improving the speed and usability of transmissions. These systems were computer-based and could consolidate, maintain, and provide on-demand the large amounts of information upon which tactical decisions are based. The Missile Minder system also incorporated computer-based logic to provide operators with recommendations concerning which assets to use to defend against what threats. Operators could thus allow the AN/TSQ-73 to automatically generate the commands necessary to conduct an air battle.

The development and fielding of the Patriot air defense missile system in the early 1980s further enhanced the capability to wage a computer-aided air battle by incorporating decision-making logic into the weapon system itself—as opposed to a separate C2 system. Decentralizing some of the engagement logic permits operators to handle a larger number of threats and speeds engagement by automating portions of the decision-making process. The utility of automating the engagement process was dramatically demonstrated with the success of the Patriot system in countering the Iraqi tactical ballistic missile (TBM) threat during Operation Desert Storm and most recently during Operation Iraqi Freedom (OIF). In both Gulf wars, TBMs were successfully engaged by Patriot employed in a fully automatic, operator-monitored mode. The down side of these successes was an unacceptable number of fratricidal engagements attributable to track misclassification problems, particularly during OIF.

In spite of these problems, Patriot as employed in the recent Gulf wars portends the future for air defense C2. The ready availability of automated support has irrevocably changed the face of air defense C2 operations. More and more of the conduct of an air battle can and will be automated. The marked increase in threat approach speeds demands that the engagement process be augmented by technology. Operators must have automatic data processing available to rapidly and simply provide the information required for engagement decision-making. The complexity of contemporary air defense C2 operations is simply too demanding to consider any other approach.

1.3 Air Defense Command and Control as Process Control

Air defense C2 systems are complex, person-machine process control facilities. Hence, to open our discussion of automation and its impact on air defense C2 operations, let us consider a notional view of the humans' role in process control. Figure 1 is a representative schematic of a process control environment containing both human and machine components. The right half of the figure represents the machine subsystem. Visual and other displays represent equipment status and the external environment in a form that human operators can understand. Controls allow the human operators to make changes in equipment status and affect the external

environment. The human subsystem is represented by the left side of figure 1. Information from the displays is first perceived. This information is processed and decisions are made. Motor responses are then made to alter control settings and impact the controlled process.




Figure 1. A representative person-machine process control system.

The vertical line separating the human subsystem from the machine subsystem represents the person-machine interface. Information flows across the interface in both directions—from the machine subsystem to the operators and from the operators to the machine subsystem. The system depicted in figure 1 is referred to as “closed loop” because one can start at any point in the diagram and work around the complete system, eventually returning to the starting point. At the concept level, the air defense C2 process also can be characterized as a closed-loop system

1.4 The Changing Role of Operators Under Automation

As machine technology has evolved, the operators’ role in air defense C2 systems also has changed. Early air defense C2 systems required operators to perform in a traditional manual control role. In later generations of systems, functions that had been in the domain of the human operators were relegated to the machine. These include not only functions that are purely psychomotor and sensory, but also those that involve cognitive skills. One way of characterizing the evolution of the operators’ role over the past several decades is that it has shifted from traditional operator to supervisory controller. The operators’ role has changed from simply

operating the system against a single threat (or against multiple threats in a sequential manner) to managing a resource pool in order to defeat an enemy air threat. Hopkin (1992) refers to this switch in emphasis as changing from a tactical to a strategic orientation.

New technologies applied to the control of complex person-machine systems can have a significant impact on operator performance and training requirements. The crux of the problem of new technologies applied to system control is that they tend to remove human operators from moment-to-moment, on-line control and relegate them to the role of supervisory controllers. A variety of research has indicated that the consequences of this change are not always positive. As discussed in the next section, the problems associated with supervisory control generally fall into one of two categories: (a) loss of situational awareness (SA) and (b) skill decay. The essential idea of SA is that operators must keep track of considerable information from a variety of sources over time and organize or interpret this information to behave appropriately (Howell, 1993). Endsley (1996) further defines SA as (a) the perception of elements in the environment, (b) the comprehension of their meaning, and (c) the projection of their status in the near future. Automation does not eliminate the requirement for operators to maintain SA, but improper implementation can make it more difficult for them to do so.

Parasuraman and Riley (1997) argue that automation does not replace human operator tasks; rather it changes the nature of the work that the operators do. It does this in ways that are often unanticipated by system designers and managers. In many process control situations—including air defense C2—we do not have a good understanding of what the supervisory control role consists of, how it should be supported, or how operator trainees should be prepared for the job. The preponderance of theory and empirical evidence suggests, however, that the job of supervisory controller is quite different from that of a traditional operator. And to maintain effectiveness, these differences must be reflected in system design, performance support features (i.e., job aids), and operator selection and training. The sections to follow address these topics as they relate to automated air defense C2 operations.

2. Automation and Process Control

2.1 An Overview of Automation

As noted previously, the next generation of air defense C2 systems will rely extensively on automated operations. A brief introduction to the topic of automation and its potential impact on air defense C2 was also provided. The objective of this section is to expand upon that earlier material and further discuss the impact of automation on process control.

Sheridan (1992, p. 3) defines automation as “the automatically controlled operation of an apparatus, a process, or a system by mechanical or electrical devices that take the place of human

organs of observation, decision, or effort.” Automated systems are typically developed for three reasons (Wickens, 1984):

1. To perform functions that humans cannot perform because of inherent limitations.
2. To perform functions that the human operators can do but perform poorly or at the cost of high workload.
3. To augment or assist performance in areas where humans show limitations.

Automation of the third kind is similar to the second, but is distinct in that automation is not intended as a replacement for core operator tasks but to assist the operator on peripheral tasks necessary to accomplish a major task. The cruise control mechanism on an automobile is a simple example of the third class of automation. A skilled driver is fully capable of manually controlling an automobile’s speed. But under certain conditions, he or she may choose to activate the cruise control mechanism and permit it to take over the speed regulation aspect of the driving task.

In another view, Bainbridge (1987) characterizes automation as “classical” versus “contemporary.” Classical automation is the older and more traditional variety and primarily involves the allocation of human physical functions to the machine. “Hard” automation is another term for classical automation (Shanit, Chang, & Salvendy 1987). Contemporary automation is information-technology-based and also involves allocating human sensory and cognitive functions to the machine. Zuboff (1988) coined the term “infomation” to refer to this second category of automation.

Zuboff uses the expression “information technology” to refer to the convergence of several streams of technical developments, including microelectronics, computer science, telecommunications, software engineering, and systems analysis. She comments that information technology has a unique capability to restructure operations that depend on information for transaction, record keeping, analysis, C2, and communications. Air defense C2 and many other contemporary process control settings are primarily concerned with information-technology-based automation. Henceforth, when the term automation is used without any modifiers, it refers to information-technology-based automation.

The popular concept of automation is that of a complex of machines performing their intended function with little or no human intervention, other than as a tender. Scenes from automated automobile factories and similar industrial facilities continually reinforce this notion. Sheridan (1992) remarks that there is an unfortunate tendency on the part of laymen to view automation as an all-or-none phenomenon. That is, a system is controlled either manually or automatically, with nothing in between. He notes, however, that all-or-none control is the exception rather than the rule. Automated systems that do not leave some residual functions for the human operators (as opposed to humans as maintainers or provisioners) are rare.

Wickens (1984) comments that automation is represented by a continuum and varies from that which totally replaces the human operator by a computer or machine to computer-driven assists that provide some unburdening of an overloaded operator. Sheridan (1992) presents the taxonomy of potential levels of information-technology-based automation shown in table 1.

Table 1. Levels of contemporary automation.

1. The computer offers no assistance; the operator must do it all.
2. The computer offers a complete set of actions, and...
3. Narrows the selection down to a few, or
4. Suggests one, and
5. Executes that suggestion if the operator approves, or
6. Allows the operator a restricted time to veto before automatic execution, or
7. Executes automatically, then necessarily informs the operator, or
8. Informs the operator only if queried, or
9. Informs the operator after execution if it, the computer, decides to.
10. The computer does everything and acts autonomously, ignoring the operator.

Source: Adapted from Sheridan (1992)

Sheridan (1992, 2002) also introduces the notion of shared versus trading of control. Shared control means that the humans and the computer control different aspects of the system at the same time. Under a shared control regimen, the computer is used to extend the humans' capabilities beyond what they could achieve alone or it can partially relieve the humans making their job easier. Trading of control refers to a situation where the computer backs up or completely replaces the humans. Backing up the humans means that the computer "picks up the slack" when the operators falter. Sheridan also comments that there are forms of cooperative control, where control is initiated by one party (humans or computer) and the other then refines it.

Sheridan remarks that the central human performance issue in shared control is, "Which tasks should be assigned to the humans and which to the computer?" In trading control, the primary issue is, "When should control be handed over to the computer and when should it be taken back?" Following Sheridan's distinction, a Patriot operator switching from semi-automatic to automatic mode for a TBM engagement is trading control to the computer.

2.2 Human Performance and Automation

As noted in the previous discussion, the dominant human performance theme in automation is function allocation between human operators and machine subsystems. For most of human history with machines, the allocation of work functions to people or machines has not been problematic. Early forms of machine aiding and automation primarily involved replacing human physical activities with machines. In most situations, the replacement of humans by machines met with little resistance. Replacement was viewed as ennobling, since humans could thus escape the physical drudgery that has been the norm throughout history.

As technology has advanced, it has become increasingly possible and even cost-effective to replace human performance of system functions with machine processing. Moreover, in advanced, capital-rich countries the allocation decision has often been simple. One merely follows Paul Fitts' (1951) dictum: Render unto the human that at which the human excels and unto the machine that which the machine excels. In recent years, however, a Pandora's Box of sorts has been opened regarding automation. Advances in information technology have, for the first time, made it possible to replace human sensory and cognitive activities by machine processes. Because of these developments, the issue of which component—human or machine—excels at what function is no longer clear-cut. In a sense, the machine has become a competitor for much of the humans' role in the person-machine system. And in some situations, it has become legitimate to ask not what the humans' role is, but whether humans are required at all, other than as a tender—a maintainer and supplier of provisions to the machine.

Since the focus of human performance concerns in automation revolve around the issue of function allocation, a brief review of function allocation concepts, methods, and problems follows. The first formal treatment of the topic is attributed to Paul Fitts (Fitts, 1951). The result of this early effort was "Fitts' List," which qualitatively describes the relatively advantages of human versus machines. Even now, Fitts' List is referenced in most introductory human factors texts. As noted previously, the basic tenet underlying Fitts' approach to function allocation is summarized as, "Render unto the human that at which the human excels, and unto the machine that at which the machine excels." People are flexible but inconsistent, whereas machines are consistent but inflexible—until recently, that is.

Fitts' approach to function allocation worked reasonably well as long as machines were not too "capable;" that is, when machine processing involved little or no "intelligence." By 1963, however, researchers such as Jordan began to criticize Fitts' so-called "engineering" approach to function allocation. Jordan suggested that the notion of human-machine comparability (the essence of Fitts' approach) be replaced by one of complementarity. Instead of casting the function allocation problem in terms of whether a task should be performed by a human or machine, it is preferable to view tasks as being performed by people and machines working together. Functions must be shared between people and machines and not just allocated to one or the other. Jordan (p. 165) also cautions system designers that, "Man is not a machine, at least not a machine like the machines that men make."

Jordan's ideas were interesting but ahead of their time. His ideas do not appear to have had much immediate impact on function allocation practice. Main-line thinking about person-machine function allocation continued along the engineering path. Chapanis (1970) is often credited with the next major installment to function allocation philosophy. His point of view is summarized as, "Let the machine do it." Under Chapanis' approach, one simply allocates to the machine anything that the machine is capable of doing. This approach is often justified for two reasons: (a) machine reliability can be increased less expensively than human reliability, and (b) there are many systems where overall system performance can be enhanced by assigning most functions to machine subsystems.

The principle flaw with the let-the-machine-do-it approach to function allocation is that humans typically are left with those functions that are too expensive or too hard to automate. In many situations, these residual functions result in a human operator task set that is termed “unreasonable” (Kantowitz & Sorkin, 1987). Unreasonable task sets can arise in two ways: First, the workload created by the task set may not match human capabilities in either the skills or the level of effort required. Although human operator overload is the first problem that designers look for in this respect, underload can be even more troublesome. Vigilance studies indicate, for example, that it is difficult even for highly motivated operators to maintain effective visual attention toward a source of information on which little is happening for more than about half an hour (Davies & Parasuraman, 1982). Designers often try to counter the vigilance problem by requiring operators to keep a log. Unfortunately, operators can make log entries without attending to their significance. Mosier and Skitka (1996) note that reliance on automated decision making can make human operators less attentive to contradictory information sources—a problem referred to as automation bias, or unwarranted over-reliance on automation.

The second way that an unreasonable task set can arise is when the functions left over for the human do not form a coherent set. If the operator cannot readily understand the logic behind a set of functions, performance will be impaired. Effective human participation in automated operations requires the development and use of a mental model of the controlled process. A mental model represents the human’s perception of the environment and concept of how system inputs translate into system outputs (Rasmussen, 1986). An incoherent set of residual functions results in a fragmented job and makes it difficult or even impossible for human operators to form a suitable mental model. Kantowitz and Sorkin (1987) remark that an incoherent task set is a more subtle and insidious problem in automated system design than is simple underload or overload.

A function allocation approach in which humans are left with whatever the machine cannot do either technically or cost-effectively is often evidence of an overall design philosophy that ignores humans. Designers who adhere to such a philosophy often would prefer not to have humans in the system at all because humans are a source of uncontrollable performance variation. The presence of humans is rationalized as a subsystem of last resort. If the machine subsystem fails, then the humans can intervene. However, if the human operator has not been left with a coherent set of functions—something reasonable to do when the system is operating normally, it is unlikely that he or she can function effectively in a manual backup mode. Situational awareness, the product of a coherent task set and meaningful participation in the controlled process, is a prerequisite for effective intervention.

System developers faced with the realities of a messy world do not design systems in the tidy manner portrayed in most textbooks. One of the better attempts to capture the spirit of real-world system design is discussed in Bailey (1983). Bailey characterizes allocation results as:

- Allocated to machine by management
- Allocated to human or machine by requirements
- Allocated by a systematic allocation procedure
- Unable to allocate

Bailey comments that in an ideal world a system designer would be given a set of system goals and could then derive the system functions needed to realize these goals. The designer would have complete freedom to allocate functions to humans or the machine. In reality, however, some of the designer's latitude is often preempted by management. Management will decide *a priori* that certain functions are to be assigned to machines. Although these decisions may not be the result of any systematic analysis, the designer is stuck with them.

System performance requirements often dictate other allocation decisions. It would be totally unacceptable, for example, to permit a computer to launch nuclear missiles without human consent. Similarly, it would be unwise to allocate complex computational or data processing activities to humans when a machine is available to perform them. Simply stated, some system performance requirements dictate the preferred function allocation scheme.

A third class of functions can be allocated either to humans or to machines. And here, the traditional function allocation methods are useful. Kantowitz and Sorkin (1987) remark that traditional methods are often more applicable when they are applied to a restricted set of system functions—the subset that remains after functions fitting into the other categories have been eliminated.

Finally, there is often a subset of functions that the designer is unable to allocate directly. In many instances, additional analyses are required to break down these functions into sub-functions that can be allocated to humans or to machines. Formal function allocation methods can then be applied at the sub-function level.

Parasuraman, Sheridan, and Wickens (2000) argue that decisions to apply automation to specific system functions should take into account the need for active human participation in the control process, even if that involvement reduces system performance from what might—theoretically—be achieved with a fully automated solution. Factors that should guide the level of automation for a particular function include (a) automation reliability, (b) the consequences of automation failure, (c) the potential impact on the operators' cognitive workload and SA, (d) the potential for operator complacency (unwarranted overtrust), and (e) the potential for and impact of operator skill degradation.

The most recent development in function allocation theory and practice is the concept of adaptive automation (see Kantowitz & Sorkin, 1987; Sheridan, 1987; Moray, 1990; Sheridan, 1992; Parasuraman, Mouloua, & Molloy, 1996; Sheridan, 2002). Under an adaptive automation scheme, an operator is able to decide function allocation on-line and often can select from

several options. The function allocation problem is, therefore, not one of allocating functions between human and machine once and for all, but dynamic allocation and re-allocation in real time as the process requires. Parasuraman et al. (1996) found that adaptive automation resulted in superior performance when compared to static allocation of a function.

In a sense, dynamic allocation merely reflects technological progress. At one time, a system's function allocation scheme was fixed in hardware. Hardware is inflexible and difficult and costly to change. As technology has advanced, however, more and more of what comprises the "system" is incorporated into software. And software may allow changes, even in real-time.

3. Human Supervisory Control

As discussed in the previous section, a key notion in the transformation of the operator function in the face of increased system complexity and automation is the idea of the operator as supervisory controller. Sheridan (1992) defines supervisory control as

...one or more operators [who are] continually programming and receiving information from a computer that interconnects through artificial effectors and sensors to the controlled process or task environment (p. 1).

In a supervisory control situation, the human operators typically handle the higher level tasks and determine the machine's goals while moment-to-moment operations are directed by the machine.

Figure 2 illustrates the basic paradigm of supervisory control (see Sheridan, 1987; Rasmussen, 1990; Sheridan, 1992, 2002). Under the supervisory control paradigm, the human operators interact directly with a dual-purpose computer, denoted the Human-Interactive Computer (HIC). One of the HIC functions is C2 and the other is to serve the operator as a decision aid. The HIC interacts with a Task-Interactive Computer (TIC), which, in turn, directs the controlled process. During normal operations, the human operators do not interact directly with the TIC.




Figure 2. Basic supervisory control paradigm.

Source: Adapted from Sheridan (1992)

The initial concepts underlying supervisory control theory grew out of Sheridan's interest in problems resulting from delayed communications in space flight (see Moray, Ferrell, & Rouse, 1990). In a long-range space flight, there is a communications delay (ranging from several minutes to several hours) between the space vehicle and the ground control station. This communications delay makes it difficult for the human ground controllers to exercise real-time C2 of the space vehicle. Sheridan proposed that a solution to the time delay problem might be to partition control intelligence between the human ground controllers and the on-board systems of the space vehicle. This suggestion led to the general notion of partitioning control intelligence in complex systems. And partitioning intelligence between human and machine components in an automated system, where the human is causally distant from the effects the machine exerts on the process, remains the central theoretical issue in supervisory control.

3.1 Traditional Versus Supervisory Control

Sheridan (2002) notes that in a traditional person-machine system, the human operator manipulates the controls of the actuators and machines directly. Furthermore, the operator can see either the results of these manipulations on the state of the process or at least view a direct representation of them by means of machine sensors and displays. Bainbridge (1987) refers to

this situation as “on-line” control. As machine capabilities have increased and particularly with the increased use of computer control, the situation vis-à-vis the operator and the controlled process has changed. With much of the control loop now assigned to machine subsystems, the humans are less tightly coupled to the machines they control. The human operators do not receive information from or exercise control over the process directly. Rather, control is exercised through a machine intermediary which controls the process and provides information to the humans (reference figure 2). It might be said that the human operators monitor and control a computer, which, in turn, directs the controlled process. Control intelligence is thus partitioned between the humans and the computer controller.

In another perspective on the issue of traditional versus supervisory control, Hopkin (1992) remarks that the role of the humans in a supervisory control situation is strategic as opposed to tactical control. Instead of focusing on and interacting intensively with individual entities—what Hopkin refers to as tactical control, the supervisory controllers use the capabilities of the computer controller to adjust global aspects of the process so that tactical control requirements are met. Under contemporary air defense C2 concepts, this shift from tactical to global control is the essence of interactive air battle management.

Traditional and supervisory control are not dichotomous operating modes. Like automation, system control modes are best characterized as arranged along a continuum ranging from traditional control through full automation with loose human oversight. Moreover, supervisory control used as a primary operating mode should not preclude operators from resorting to traditional, tactical control in specific situations. Effective system control often requires operators to use a combination of the two modes, and the system must permit both modes to be applied. The point to emphasize is that with most new air defense C2 systems the dominant control mode is supervisory. Supervisory control is different conceptually from traditional control. Moreover, supervisory control imposes different system design, human performance, and training requirements than is the case with traditional control.

3.2 The Human Operator’s Role as a Supervisory Controller

When addressing the human operators’ role in a supervisory control setting, it is instructive to begin with a discussion of Rasmussen’s (1986) supervisory control taxonomy. This taxonomy is shown graphically in figure 3. Under Rasmussen’s taxonomy, human tasks in a control system can be classified into one of three categories, denoted skill-based behavior (SBB), rule-based behavior (RBB), and knowledge-based behavior (KBB). Skill-based behaviors consist of sensory and motor performances during acts that, after a statement of intent, take place without conscious control as smooth, automated, and highly integrated behaviors. A simple example of SBB is entering commands into a C2 computer.




Figure 3. Rasmussen's supervisory control taxonomy.

Source: Adapted from Rasmussen (1986)

In RBB, the task sequence is goal-oriented and consciously controlled by a stored rule. This governing rule may have been (a) derived empirically during previous operations, (b) communicated from another person's know-how, or (c) prepared on occasion through conscious problem solving and planning. The boundary between SBB and RBB is not distinct; it depends on both the level of training and attention of the operator. Conscious RBB for an inexperienced operator might be automatic SBB for a more experienced one.

When the operator is faced with situation for which no explicit rules are available, performance control shifts to a higher conceptual level in which actions are goal-oriented and determined on occasion through conscious problem solving and planning. Rasmussen refers to this later category of human performance as KBB. The structure of KBB is a function of the operator's skill level, experience, and comprehension of the tactical situation.

Rasmussen's supervisory control taxonomy provides a useful perspective on the human performance requirements underlying supervisory versus traditional control. Simply stated, a supervisory control regimen emphasizes and retains operator decision-making and problem-solving tasks while relegating most direct sensory and psychomotor tasks (SBB) and many rule-based performances (RBB) to machine subsystems. Activities in the skill-based performance

domain can be allocated either to humans or to the machine. A good rule is to make any skill-based activities allocated to humans as simple and error proof as possible (e.g., a point-and-click operation with a mouse as opposed to a complex series of keystrokes).

It might be said that, by definition, the knowledge-based performance domain is the exclusive preserve of the human operators. Parasuraman and Riley (1997) note that automation of knowledge-based functions such as decision making, planning, and creative thinking “remains rare.” These authors go on to state (p. 232) that “despite more than three decades of research on AI (artificial intelligence), neural networks, etc., enduring transfer of thinking skills to machines has proven very difficult.”

Allocating the so-called intelligent aspects of system control between humans and the machine involves partitioning the rule-based performance domain into a subset assigned to the machine and a second subset assigned to the human operators. Following this logic, adaptive automation involves partitioning and re-partitioning the set of rule-based performances between humans and the machine subsystem. Figure 4 illustrates the relationship between Rasmussen’s supervisory control taxonomy and person-machine function allocation in automated systems.




Figure 4. Rasmussen’s supervisory control taxonomy and human-system function allocation in an automated system.

In another characterization of the operator’s role in supervisory control, Sheridan (1992, 2002) notes that operators typically retain a set of generic residual functions. He adds that these functions are consistent across a range of systems. Sheridan refers to these residual operator activities as roles, but we prefer to use the term functions. Sheridan’s residual operator functions are listed as follows:

1. Monitoring automatic control. Allocate attention among displays to ensure that the system is operating well and close to setpoint.
2. Intervening to update instructions or assume direct control. Interrupt operations to send commands to the computer when abnormal conditions occur or the machine drifts from setpoint. Assume direct control of the TIC or the controlled process when the situation requires such action.
3. Learning from experience. Benefit from experience to develop an increasingly efficient mental model of the controlled process which can be used to predict process behavior in situations not yet experienced.
4. Teaching the computer. Instruct the machine to change its mode of operation, reset control parameters, modify software, etc.
5. Planning. Choose strategic or tactical options for achieving the system's overall goals.

Sheridan describes these generic operator functions as organized in three nested loops, as illustrated in figure 5. The inner loop, Monitoring, closes on itself. Evidence of something out of the ordinary or completion of one part of the Monitoring cycle leads to additional investigation and monitoring. The middle loop closes from Intervention back to Teaching. Human intervention usually leads to the specification and programming of a new goal state for the process. Finally, the outer loop closes from Learning back to Planning. Intelligent planning for the next event is usually not possible without learning from the last one.




Figure 5. Functional and temporal nesting of generic supervisory control functions.

Source: Sheridan (1992)

The progression of loops depicted in figure 5 also suggests a natural learning sequence for supervisory control. In order, the elements of this sequence are: (a) Monitor and Intervene, (b) Learn and Teach, and (c) Plan. Also note that this training sequence would best be implemented over the course of an operator's career and should be interspersed with periods of on-the-job experience in the role segment for which training has just been completed.

Prior to concluding the present section, it is instructive to comment on two additional aspects of operator performance in a supervisory control environment. First, recall that the HIC is a dual purpose computer: C2 and decision aiding. Thus far the discussion has addressed C2, but has said little about decision aiding. Figure 6 illustrates the ideal types of decision and job aiding provided by a computer aid for each of the three levels of Rasmussen's taxonomy. Note the differences in the types of queries input by the operator for each level of taxonomy and the nature of the information provided by the computer decision aid.




Figure 6. Operator interaction with a computer-based decision aid.

Source: Adapted from Rasmussen (1986)

As shown in figure 6, skill-based job aiding consists of control command clarifications and demonstrations. A common example of skill-based aiding is the "Help" function provided with many software packages. Users are able to obtain on-line a function-by-function clarification of the package's command structure or even a short tutorial (in essence, a demonstration) merely by invoking the Help option. Moreover, the Help function can be activated without leaving the application at hand.

Rule-based job aiding consists of providing operators with the capability to peruse and receive information about the rule structure governing a particular application. The technology of practical rule-based job aiding in real-time is still in its infancy. However, several operational rule-based job aids are in current use. A good example of an operational rule-based performance aid is the Air Force's Observing System for Critique, Advice, and Review (OSCAR) (Obermayer, 1991). OSCAR is a PC-based capability intended to critique pilot performance on an F-16 Air Intercept Trainer. The system is able to parse the continuous record of a simulated air engagement and provide trainees with a maneuver-by-maneuver evaluation of their performance vis-à-vis the tactics manual. If requested, the rule structure governing the critique can also be displayed, along with specific references from the tactics manual or other performance guides.

The final aspect of supervisory control performance support concerns knowledge-based job aiding. For most applications, the first step in defining knowledge-based performance aids is to proceduralize the knowledge-based aspects of the situation. Proceduralization involves generally specifying the activities to be performed and then identifying the solution approach to be employed. Spreadsheets and other management support tools are common examples of knowledge-based job aids. The computer-based mission planning aids that are routinely used to support air operations and other military activities are yet another example of practical knowledge-based job aiding.

Note that there is a qualitative difference between skill- and rule-based job aiding versus knowledge-based aiding. Skill- and rule-based aids provide circumscribed guidance in response to a specific request: How do I? What is the governing rule for this situation? Knowledge-based job aids, on the other hand, often consist of a tool (spreadsheet program, planning aid, etc.) that permits the operator to use his or her job knowledge and performance goals to formulate a solution to a general class of problems.

The second comment concerns display format adaptivity in a supervisory control environment. Adaptive displays are variable in format or logic as a function of mission stage or operating conditions. Sheridan (1987) remarks that operator displays in most current industrial and aerospace systems have fixed formats. Future control stations may, however, require adaptive displays. He further notes that little concrete work on the development, use, or effectiveness of adaptive displays has yet been done.

Sheridan (1987) provides a useful historical-technical perspective on supervisory control. To begin, he positions various work roles in a two-dimensional Cartesian space, with the axes defined by (a) the degree of automation (Manual through Fully Automated) and (b) the complexity and unpredictability of the job situation (referred to as Task Entropy). Entropy is a measure of the level of uncertainty in the operating environment. An example of the resulting space is shown in figure 7.




Figure 7. Potential combinations of human and computer control.

Source: Adapted from Sheridan (1992)

Sheridan argues that supervisory control is properly viewed as a frontier line advancing gradually toward the upper right hand corner of the space—toward a Perfect Robot. The impetus for the advancing frontier line is provided by developments in information, sensor, and display technologies. Current machines are capable of performing tasks that several years ago would have been the sole responsibility of human operators. And tasks that presently are beyond the capabilities of machines will in the future be routinely handled by computers.

One implication of Sheridan's observation is that it may prove difficult to derive a general theory of supervisory control. Moray (1990) alludes to this possibility with his observation that the problem of operator intervention in automated system processing remains without a solid theoretical foundation. He adds that there is no adequate theory or model for function allocation in automated systems. Moray notes that various approaches to function allocation (e.g., Fitts' List and follow-on adaptations) have rapidly become out of date as technology has advanced. In Moray's view, the solution is not one of allocating functions between the human and the machine once and for all, but adaptive automation as the control situation demands

Recent developments in automation theory and practice tend to support Moray's view. Within the domain of air traffic control, for example, Hopkin (1992) comments that progressive automated assistance is being pursued as an operational concept for future automated control systems. Under this concept, the operator can select from among several pre-programmed levels and types of automated assistance. In air defense C2, the Patriot system employs a limited form of progressive automated assistance with its semi-automatic and automatic modes of operation. When the operator shifts from semi-automatic to automatic mode, he or she, in essence, changes the person-machine function boundary in real time. Activities that had been the operator's responsibility are now shifted to the computer and can be taken back.

4. The Path Forward

Information technology is driving the operator's role in complex, computer-aided human-machine systems increasingly toward that of a supervisor of automated processes and an intervener in the control process for abnormal situations. This shift is not made at the discretion of system designers. Designers are not free to opt for or against casting the operator as a supervisory controller. Like it or not, the role shift is a consequence of the change from manual to automated process control.

A number of views of the problems of automation and supervisory control are discussed in the present paper. There is, however, little agreement that any of the current concepts provide a framework for supervisory control having sufficient generality or robustness to serve as a basis for system design or for operator selection and training. Nonetheless, there are several consistent observations that can serve as guidelines for system development. The most significant of these observations are addressed in the paragraphs to follow. We begin our discussion with a review of problems associated with undisciplined automation. The term undisciplined automation refers to automated systems developed without regard for proper human-system integration.

4.1 Problems with Undisciplined Automation

There is a tendency among system designers with little training in human performance theory and methodology to assume that automation is innately beneficial. Research in a number of areas suggests, however, that such is not always the case. The paragraphs to follow highlight and discuss problems that frequently occur when automated systems are developed without regard for the human component.

Newly automated systems seldom provide all anticipated benefits. In a survey of user experiences with automated systems, Boehm-Davis, Curry, Weiner, and Harrison (1983) found that automated systems rarely live up to their initial billing. These authors also comment that first-time users of automated systems must anticipate a debugging period during which the

system's actual capabilities and limitations are determined. It is often necessary for field users to determine how they should practically employ the system, as opposed to unquestioningly using it the way that system developers think it should be used. System developers often fail to anticipate problems that an automated system will create. Unquestioning acceptance of an automated system opens the door to what Parasuraman and Riley (1997) refer to as "automation misuse," or over-reliance on automation—which can result in monitoring failures and decision biases.

Increased monitoring load. Sheridan (1987) notes that automation may change the nature of the operator's job, but does not always simplify it. Automated systems often are characterized by a proliferation of components brought on by increased system complexity. Under an automation regimen, the operator often has less to do moment-to-moment, but as a consequence of an increased number of components, he or she has more indications of system status to monitor.

False sense of security. A total belief in the system's infallibility (i.e., it's always right) can lull the operator into a false sense of security with the result that he or she will not make checks that would otherwise be advisable (Mosier & Skitka, 1996). Kantowitz and Sorkin (1987) comment on the "flip" side of this problem. Frequent failures of automated equipment can lead to problems of trust in automation. And if operators have a choice in automation use, trust determines usage (Lee & Moray, 1992). Muir (1988) argues that trust in automation is affected by the same factors that influence trust between individuals, namely effectiveness and reliability. Lee and Moray (1992) present results showing that trust is slow to recover following automation failure. Situations involving systematic deficiencies result in loss of trust until operators understand fault patterns and learn to accommodate them.

Automation makes operators into system monitors rather than active controllers. Rasmussen (1986) comments that automation does not remove humans from the system. Rather, it moves human operators from direct, on-line control of system operations to higher-level supervisory control tasks. Problems arise when the automatic control system has been built because it presumably can do the job better than a human controller, but the operator is left in to "monitor" that the automated system is performing correctly (Bainbridge, 1987). There are two potential difficulties with this situation. First, the human monitor must know what the correct behavior of the system should be. Such knowledge requires a high level of system expertise and displays attuned to supervisory control requirements. Second, if system operations can be completely specified, then a computer can implement control actions quicker and more reliably than a human operator. If such is the case, then why leave humans in the control loop at all? Human operators provide little value added in a completely determined task setting. However, the more unpredictable and ambiguous the task setting, the more potential value the human operators can bring to the control situation.

Out-of-the-loop familiarity. When an operator is replaced by an automatic controller, his or her level of interaction or familiarity with the system is reduced. There is evidence that when an abnormal situation does occur, the operator will be slower to detect it and will take a longer time to “jump into the loop” and make the appropriate control action (OTA, 1984). This problem is sometimes referred to as loss of SA.

There also appear to be longer-term consequences of being removed from the control loop. Controllers may lose proficiency as they receive less and less hands-on experience. This latter situation has been identified as a significant problem for pilots who rely excessively on automated flight control systems. The remedy, of course, is regular simulator sessions covering those aspects of the job that are no longer routinely performed.

Higher-level operator error. Automation does not eliminate the possibility for human error; it merely relocates the sources of human error to a different level. Furthermore, the consequences of such error can be more significant than would be the case with a less automated system. The cascading consequences of operator errors in the Chernobyl or Three Mile Island nuclear reactor incidents represent a good case in point. Sheridan (1987) likens the potential for error propagation in an automated system to what happened to the Sorcerer’s Apprentice in Walt Disney’s Fantasia—the proliferation of “automated” brooms. He stresses that automated systems tend to be less forgiving than their less sophisticated predecessors.

Bainbridge (1987) discusses another aspect of automation and human error that is worthy of special comment. She remarks that automatic control systems can mask system failure by controlling against indicator variable changes so that a drift to failure is not apparent until it is too late to take corrective action. This problem implies that automatic control systems must monitor unusual movements of critical system variables and alert the operator to emerging trends. Bainbridge notes that a graceful degradation of performance is often cited as an advantage of humans over machines. She comments that graceful degradation is not an aspect of human performance that should be emulated in an automated control system because it complicates the task of monitoring for failure. Bainbridge argues that automatic systems should unambiguously signal their drift toward failure.

Loss of cooperation among human operators. In systems requiring crew coordination, automation tends to lessen the frequency of interaction among team members. The result is that when effective team interaction is required, coordination skills may have decayed to the point where they can no longer be performed. Moreover, because of the changed nature of the operators’ roles, automation often imposes new coordination demands (Parasuraman, Sheridan, & Wickens, 2000). In a sense, this problem is an interpersonal aspect of the out-of-the-loop familiarity problem.

Increased training requirements. Contrary to popular belief, automation does not always lessen operator training requirements. It frequently increases operator training requirements and often raises aptitude levels (OTA, 1984). In a study of manufacturing automation, the Office of Technology Assessment (OTA) also concluded that automation often brings with it a greater

need for conceptual skills on the part of operators. Blumberg and Alber (1982) also found that automation often requires a movement away from specialized job classifications in the direction of more general and highly skilled jobs.

Boehm-Davis, et al. (1983) comment that the increased training burden associated with automated systems often stems from a requirement to be able to operate the equipment in two modes, automatic and manual. Furthermore, it is often the case that automatic operation deprives the operators of practice in the manual mode. And the resulting loss of proficiency can result in performance problems when the manual mode must be activated.

Automated systems also tend to be more complex than their non-automated predecessors. This increased complexity can make system operations more difficult to learn. Moreover, if the system lacks a coherent job structure, that fact can compound learning problems associated with system complexity. From the perspective of both training and operational effectiveness, the worst possible situation is a complex automated system that does not leave the operator with a coherent job.

We remarked previously on the need for frequent simulator sessions to combat problems with out-of-the-loop familiarity. There are, however, several inherent problems with the use of simulators to maintain supervisory control proficiency. Perhaps the most serious of these problems is the difficulty of training for extreme situations. Unknown system faults cannot be simulated, and system behavior may not be known for faults that can be predicted but have not been experienced. Rasmussen (1986) comments that the skills required for performance during extreme situations are not always developed or maintained during normal training or operating modes.

Lee and See (2004) argue that automation should be designed for appropriate as opposed to greater trust. These authors assert that training is an important consideration in developing appropriate trust. They go on to state that operator training should emphasize (a) expected system reliability, (b) the mechanisms underlying potential reliability problems, and (c) how usage situations interact with the automation's technical characteristics to affect its capability. Rovira, McGarry, and Parasuraman (in press) present results illustrating one of the true "ironies of automation": The more reliable the automation, the greater its detrimental effect when it does fail.

As a final comment on training and automation, consider Bainbridge's (1987) observation that with the present generation of automated systems we may be "riding on the skills" of former manual operators. She notes that future generations of operators cannot be expected to have these skills and expresses some concern regarding where future controllers of complex systems will acquire the necessary skills. Bainbridge asks, somewhat rhetorically, whether it is possible for an unskilled operator to effectively monitor a complex system. To bring this point home, think about the situation that would result if airline pilots were not trained to fly manually, but were only taught to control aircraft using automated flight control equipment. What would happen if the automatic control equipment were to fail and require pilot intervention—during landing, say?

4.2 Developing Effective Automated Systems

In spite of the heuristic nature of many of the conclusions and observations concerning automation and supervisory control (i.e., rules of thumb as opposed to hard-and-fast rules), there are several points of general agreement to consider in the development of effective automated systems. Several of the more significant of these guidelines are discussed as follows.

Function analysis and allocation. Contrary to much current practice, automated systems require more attention to explicit human-system integration than is the case with less complex systems. As long as humans are required, designers must keep in mind that an automated system includes both people and machine subsystems. Care must be taken that a particular automation concept neither overloads nor underloads operators and that the residual operator task set constitutes a coherent set of activities. This requirement makes it imperative that early-on usability assessments of automated systems be carefully and professionally performed. The consequences of deploying an automated system having a defective person-machine integration concept can be more serious than is the case with less sophisticated systems. In less complex systems, operators are able to compensate for design deficiencies; in an automated system, operator compensation can be more difficult. In addition to careful design, both designers and users must anticipate an explicit debugging period during which the developer's operational concept is evaluated and adjusted to reflect field usage conditions.

Flexible and adaptive implementation. Weiner and Curry (1980) note that desires and needs for automation will vary across individuals and situations. Hence, automation should be flexibly available and not mandatory. These authors also remark that automated decision support systems should provide guidelines and suggestions, not commands. An accepting, forgiving system will be better accepted by users than an autocratic one.

A forgiving, fault-tolerant computer system. The computer interacting with human operators—sometimes referred to as the Human Interactive Computer, or HIC—should be designed to permit common human errors (i.e., “slips”) to occur without greatly disrupting system performance. Several recent sources in the area of software user interface design (e.g., Mayhew, 1992) provide excellent discussions of fault-tolerant person-computer interface design.

Feedback. Clear, unambiguous feedback regarding the consequences of control actions must be available to the operator at all times. Examples of such feedback include the “Are You Sure?” query sometimes encountered in commercial computer software. Furthermore, the time delay between control input and system acknowledgment or reaction must be as short as possible. Also, the issue of “Who’s in Charge” must be clearly spelled out. It must be clear to the operators who—humans or the computer—has the responsibility for a particular action.

Understanding the processing of automated equipment. Weiner and Curry (1980) argue that the operator’s understanding of the process by which the automated subsystem does its job is a major determinant of how well he or she trusts the component. If the automated subsystem is

viewed as a “black box,” it will be less trusted than if the operator understands the logic followed by the automated subsystem. These authors suggest that the automated component should be designed to perform a task in a manner that is consistent with how the user would perform it. In this manner, the user’s mental model of the automated subsystem’s performance will be consistent with the system’s actual processes.

Weiner and Curry further suggest that extensive training in the use of the automated system be provided so that (a) users appreciate its potential benefits and (b) failures and malfunctions are more readily identifiable. Of necessity, such training will often consist of training in manual operations followed by training in automated operations. If operators are to appreciate and accept the benefits of automated operations, they must have some basis for comparing the two modes of operation. These authors also argue that if the operators are to make judgments concerning the adequacy of the automated subsystem’s performance, they must have some basis for making those judgments.

4.3 Closing Perspectives: Beyond Human-Centered Design

Automation can be a two-edged sword. In response to technological and mission changes, the role of air defense C2 crew members has evolved from traditional operators to supervisors of automated processes. But as long as human control intelligence is required, automation brings with it the problems associated with supervisory control, primarily vigilance decrements and loss of SA. The irony of the situation is that by automating one does not eliminate the difficulties encountered with either traditional on-line control or classical automation. Research and experience indicate that effective supervisory control requires a skilled operator in somewhat continuous and meaningful interaction with the system. This conclusion does, however, seemingly contradict the objectives of automation, which often are to remove the operator from moment-to-moment interaction with the system.

Beyond proper human-centered design, another way to lessen the likelihood of vigilance and SA problems in supervisory control is to increase the skill level of operators. Research from performance domains similar in concept to air defense C2 indicates that experienced job performers do not need to stay as tightly coupled to the problem-solving situation as less skilled operators in order to perform effectively (Ericsson & Charness, 1994). Highly skilled job performers are able to maintain SA with less direct effort and are also able to reestablish SA after an interruption in less time and with less apparent effort.

If one accepts the notion that increasing the skill level of operators is a practical way to address some of the human performance problems associated with supervisory control, a subsequent issue concerns how sufficient levels of skill are to be developed. Chase and Simon (1973) are adamant on this point: “Practice is the major independent variable in skill acquisition” (p. 279). Ericsson and Charness (1994) assert that skilled performance is not an automatic consequence of simply having more experience with an activity. Rather, skilled performance is developed through what is termed “deliberate practice.” Ericsson, Krampe, and Tesch-Romer (1993) define

deliberate practice as “a long period of active learning during which they [job performers] refine and improve their skills under the supervision of a teacher or coach” (p. 368). Ericsson, et al.’s (1993) notion of deliberate practice is consistent with the contemporary definitions of training within the Department of Defense as “relevant practice with feedback” (Defense Science Board, 2003, p. 20).

Structuring relevant practice with feedback for a complex supervisory control situation can, however, be a non-trivial undertaking and will require many changes from traditional air defense training and staffing practices. The Boards of Inquiry (BOIs) examining the root causes of Patriot fratricide incidents during Operation Iraqi Freedom concluded that the training provided to van crews was a contributing factor. Specifically, the air defense community was criticized for training practices that “emphasize rote drills rather than high-level judgment.” What this means is that much pre-OIF Patriot training was reduced to a stimulus-response exercise with little intervening thought or judgment: If you see X...Then do Y. A rote, crew-drill approach to training might be appropriate for many aspects of air defense operations (e.g., march order, emplacement, system set-up, etc.), but it is not suitable for air battle operations or management. These require a focus on adaptive decision making within a complex and dynamic tactical setting.

The challenge facing the air defense community is to change from an emphasis on rote drills to one fostering adaptive decision making. This will require approaching van crew and battle staff training and development using a framework that has proven useful in performance domains similar to air defense. For example, McPherson and Kernodle (2003) argue that training for a complex decision-making setting must address three components of effective performance:

1. Declarative knowledge: Facts and rules governing the problem situation.
2. Tactical knowledge: Analysis, planning, self-monitoring, anticipation, etc. What to do?
3. Procedural knowledge: How to execute the intent decided in step 2.

From an air defense perspective, the first component concerns the technical and tactical knowledge underlying effective air battle operations and management. The second component is concerned with using technical and tactical knowledge of the operational setting and the system’s current status to determine an appropriate course of action (What to do). Component three then involves the button and function key (i.e., “switchology”) aspects of system operations (How to execute that action).

Our observation of training for air battle operations and management suggests that current instruction emphasizes components one and three, but gives short shrift to component two. That is, trainees are (a) given blocks of classroom instruction on air defense and system-specific knowledge and tactics and are (b) provided with follow-on simulator training addressing system switchology. The missing or poorly implemented component involves translating background knowledge and the current tactical situation into an appropriate course of action. With a missing or poorly conducted tactical component, training is effectively reduced to a series of rote drills. This is the essence of the BOI criticism of the pre-OIF training provided to Patriot crews.

The path forward in this respect is direct but not simple. In effect, the air defense community must creatively put in place the conditions for deliberate practice. These are, (a) valid training scenarios, (b) effective after-action reviews (AARs)—feedback, and (c) time to develop the necessary levels of expertise. Prospective van crews and battle staff members must be able to translate their air defense and system knowledge into an appropriate tactical plan and then know how to carry out this plan using the system as a tool.

5. References

Bailey, R. W. *Human performance engineering*. Englewood Cliffs, NJ: Prentice Hall, 1983.

Bainbridge, L. Ironies of automation. In J. Rasmussen, K. Duncan, and J. Leplat (Eds.), *New technology and human error*. New York: Wiley, 1987.

Blumberg, J.; Alber, A. The human element: Its impact on the productivity of advanced batch manufacturing systems. *Journal of Manufacturing Systems* **1982**, *1*, 43–52.

Boehm-Davis, D. A.; Curry, R. E.; Weiner, E. L.; Harrison, R. L. Human factors of flight-deck automation: Report on a NASA-industry workshop. *Ergonomics* **1983**, *26*, 953–961.

Chapanis, A. Human factors in systems engineering. In K. B. DeGreene (Ed.), *Systems psychology*. New York: McGraw-Hill, 1970.

Chase, W. G.; Simon, H. A. The mind's eye in chess. In W. G. Chase (Ed.), *Visual information processing*. New York: Academic Press, 1973.

Davies, D.; Parasuraman, R. *The psychology of vigilance*. New York: Academic Press, 1982.

Defense Science Board. *Training for future conflicts*. (Final Report of the DSB Task Force on Training for Future Conflicts). Washington, DC: Office of the Undersecretary of Defense for Acquisition, Technology, and Logistics, 2003.

Ericsson, K. A.; Charness, N. Expert performance: Its structure and acquisition. *American Psychologist* **1994**, *49* (8), 725–747.

Ericsson, K. A.; Krampe, R. T.; Tesch-Romer, C. The role of deliberate practice in the acquisition of expert performance. *Psychological Review* **1993**, *100* (3), 363–406.

Fitts, P. M. *Human engineering for an effective air navigation and traffic control system*. Washington, DC: National Research Council, 1951.

Endsley, M. Automation and situation awareness. In R. Parasuraman & M. Mouloua (Eds.), *Automation and human performance: Theory and applications*. Hillsdale, NJ: Erlbaum, 1996.

Hopkin, V. D. Human factors issues in air traffic control. *Human Factors Society Bulletin* **1992**, *35* (6), 1–4.

Howell, W. C. Engineering psychology in a changing world. *Annual Review of Psychology* **1993**, *44*, 231–63.

Jordan, N. Allocation of functions between man and machines in automated systems. *Journal of Applied Psychology* **1963**, *47*, 161–165.

Kantowitz, B. H.; Sorkin, R. D. Allocation of functions. In G. Salvendy (Ed.), *Handbook of human factors*. New York: Wiley, 1987.

Lee, J. D.; Moray, N. Trust, control strategies, and allocation of function in human-machine systems. *Ergonomics* **1992**, *35*, 1243–1270.

Lee, J. D.; See, K. A. Trust in automation: Designing for appropriate reliance. *Human Factors* **2004**, *46* (1), 50–80.

Mayhew, D. J. *Principles and guidelines in software user interface design*. Englewood Cliffs, NJ: Prentice Hall, 1992.

McPherson, S. L.; Kernodle, M. W. Tactics, the neglected attribute of expertise. In J. Starkes & K. Ericsson (Eds.), *Expert performance in sports: Advances in research on sports expertise*. Champaign, IL: Human Kinetics, 2003.

Mosier, K. L.; Skitka, L. J. Automation use and automation bias. In *Proceedings of the Human Factors and Ergonomics Society 40th Annual Meeting* (pp. 204-208). Santa Monica, CA: Human Factors and Ergonomics Society, 1996.

Moray, N. Mental models of complex systems. In N. Moray, W. R. Ferrell, and W. B. Rouse (Eds.), *Robotics, control, and society*. New York: Taylor & Francis, 1990.

Moray, N.; Ferrell, W. R.; Rouse, W. B. *Robotics, control, and society*. New York: Taylor & Francis, 1990.

Muir, B. M. Trust between humans and machines, and the design of decision aids. In E. Hollnagel, G. Mancini, & D. D. Woods (Eds.), *Cognitive engineering in complex dynamic systems*. London: Academic Press, 1988.

Obermayer, R. W. *Observing system for advice, critique, and review (OSCAR)*; Technical Report; U.S. Air Force Human Resources Laboratory: Williams AFB, AZ, 1991.

Office of Technology Assessment. *Computerized manufacturing automation: Employment, education, and the workplace*. (OTA-CIT-235). Washington, DC: US Government Printing Office, 1984.

Parasuraman, R.; Mouloua, M.; Molloy, R. Effects of adaptive task allocation on monitoring of automated systems. *Human Factors* **1996**, *38* (4), 665–679.

Parasuraman, R.; Riley, V. Humans and automation: Use, misuse, disuse, abuse. *Human Factors* **1997**, *39* (2), 230–252.

Parasuraman, R.; Sheridan, T. B.; Wickens, C. D. A model for types and levels of human interaction with automation. *IEEE Transactions on systems, Man and Cybernetics—Part A: systems and Humans* **2000**, *30* (3), 286–297.

Rasmussen, J. *Information processing and human-machine interaction: An approach to cognitive engineering*. New York: North Holland, 1986.

Rasmussen, J. Supervisory control and risk management. In N. Moray, W. R. Ferrell, and W. B. Rouse (Eds.), *Robotics, control, and society*. New York: Taylor & Francis 1990.

Rovira, E.; McGarry, K.; Parasuraman, R. Effects of imperfect automation on decision making in a simulated command and control task. *Human Factors*, (in press).

Shanit, J.; Chang, T-C; Salvendy, G. Technical and human aspects of computer-aided manufacturing. In G. Salvendy (Ed.), *Handbook of human factors*. New York: Wiley, 1987.

Sheridan, T. B. Supervisory control. In G. Salvendy (Ed.), *Handbook of human factors*. New York: Wiley, 1987.

Sheridan, T. B. *Telerobotics, automation, and human supervisory control*. Cambridge, MA: MIT Press, 1992.

Sheridan, T. B. *Human and automation: System design and research issues*. New York: Wiley, 2002.

Weiner, E. L.; Curry, R. E. Flight deck automation: Promises and problems. *Ergonomics* **1980**, *23*, 995–1012.

Wickens, C. D. *Engineering psychology and human performance*. Columbus, OH: Merrill, 1984.

Wickens, C. D. Information processing, decision-making, and cognition. In G. Salvendy (Ed.), *Handbook of human factors*. New York: Wiley, 1987.

Zuboff, S. *In the age of the smart machine: The future of work and power*. New York: Basic Books, 1988.

Distribution List

ADMNSTR
DEFNS TECHL INFO CTR
ATTN DTIC-OCP (ELECTRONIC COPY)
8725 JOHN J KINGMAN RD STE 0944
FT BELVOIR VA 22060-6218

DIRECTORATE OF COMBAT
DEVELOPMENT
ATTN A MARES (2 COPIES)
5800 CARTER RD
FT BLISS TX 79916

DIRECTOR
DIRECTORATE OF COMBAT
DEVELOPMENT
ATTN COL D KIRBY
5800 CARTER RD
FT BLISS TX 79916-7001

DIRECTOR
DIRECTORATE OF TRAINING, DOCTRINE,
AND LEADER DEVELOPMENT
ATTN COL R BURNS
2 SHERIDAN RD, BLDG 2
FT BLISS TX 79916-7001

OFFICE, CHIEF OF AIR DEFNS ARTILLERY
ATTN LTC R STARKEY
2 SHERIDAN RD, BLDG 2
FT BLISS TX 79916-7001

TRADOC SYSTEM MANAGER-LOWER
TIER
ATTN COL R JASSEY
BLDG 12, PERSHING RD
FT BLISS TX 79916-7001

TRADOC SYSTEM MANAGER-UPPER TIER
ATTN COL S PETERS
BLDG 12, PERSHING RD
FT BLISS TX 79916-7001

COMMANDING GENERAL
US ARMY AIR DEFNS ARTILLERY CTR
AND FT BLISS
ATTN MG M VANE
FT BLISS TX 79916-7001

US ARMY RSRCH LAB
ATTN AMSRL-ARL-CI-OK (TECH LIB)
BLDG 4600
ABERDEEN PROVING GROUND MD 21005-
5066

US ARMY RSRCH LAB
ATTN AMSRL-ARL-HR-M F PARAGALLO
ATTN AMSRL-ARL-HR-MB
C GIAMMANCO (4 COPIES)
ATTN AMSRL-ARL-HR-MB J HAWLEY
(5 COPIES)
BLDG 459
ABERDEEN PROVING GROUND MD 21005-
5066

US ARMY RSRCH LAB
ATTN AMSRD-ARL-CI-OK-T TECHL PUB
(2 COPIES)
ATTN AMSRD-ARL-CI-OK-TL TECHL LIB
(2 COPIES)
ATTN AMSRD-ARL-D J M MILLER
ATTN IMNE-ALC-IMS
MAIL & RECORDS MGMT
ADELPHI MD 20783-1197